

GRATIA STORAGE ACCOUNTING

Project History

2

- Grid Accounting (Gratia) system was created for batch systems and Linux process accounting
- Designed and developed by a group of developers, lead by Philippe Canal:
 - Chris Green
 - Brian Bockelman
 - Karthik Arunachalam and others
- In production since 2006 at FNAL
- Was adopted by the Open Science Grid grid-wide accounting system in 2007
- I am taking over the Project Manager responsibilities starting October 1st, 2011

Brief Overview

3

- Information is generated by various probes and sent to Gratia collectors via Gratia API
- Collects information about:
 - ▣ Batch and glide-in jobs (condor, lsf, pbs)
 - ▣ Linux process accounting
 - ▣ Various Metrics (RSV probes)
 - ▣ File transfers
 - ▣ Storage Usage
- Supports multiple collectors
- Supports hierarchical forwarding between collectors
- Allows data filtering and replication
- Provides means to generate various reports

Storage Accounting

4

- In 2009, USCMS requested that Gratia provide the means to collect and store storage utilization data
- Two new accounting entities have been added:
 - Storage Element
 - Storage Element Record
- StorageElement is used to describe static information and storage topology (name and type of storage, storage area parent, etc)
- StorageElementRecord is used to store dynamic information: space measurement
- The design is based on OGF Usage Record standard
- Storage gratia probes are currently developed for:
 - dCache
 - HDFS
 - Xrootd

StorageElement Table

5

- UniqueID (Site Name: Storage Name)
 - FNAL_ITB:Area:public-link-group
 - Nebraska:Quota:/user/engage
 - FNAL_PUBLIC_DCACHE:Pool:w-stkendca20a-2
- Grid
- Site Name
- Storage Element Name
- ParentID (SE Unique ID: eg FNAL_GRIDWORK)
- SpaceType
 - SE
 - Area
 - Pool
 - Quota
- Implementation
 - dCache
 - Hadoop
- Version
- Status
 - Production
 - Closed
- VO
- OwnerDN
- Timestamp
- ProbeName

StorageElementRecord Table

6

- UniqueID
- RecordMeasurementType
 - Logical
 - Raw
- StorageType
 - disk
- TotalSpace
- FreeSpace
- UsedSpace
- FileCountLimit
- FileCount
- ProbeName
- Timestamp

dCache Storage Probe

7

- Gets Information from dCache Information Provider
- Runs as a cron job
- Reports information about
 - Storage Element
 - Group links
 - Pools (optional)
 - Space Tokens
- Examples (xml message sent by a probe)

```
<?xml version="1.0" encoding="utf-8"?><StorageElementRecord xmlns:urwg="http://www.gridforum.org/2003/ur-wg">
<RecordIdentity urwg:createTime="2011-09-01T13:24:40Z" urwg:recordId="fermicloud042.fnal.gov:7287.0"/>
  <UniqueID>FNAL_PUBLIC_DCACHE:SE:FNAL_PUBLIC_DCACHE</UniqueID>
  <MeasurementType>raw</MeasurementType>
  <StorageType>disk</StorageType>
  <TotalSpace>97375107835914</TotalSpace>
  <FreeSpace>9907103813958</FreeSpace>
  <UsedSpace>87468004021956</UsedSpace>
  <Timestamp>2011-09-01T13:24:39Z</Timestamp>
  <ProbeName>dCache-storage:fermicloud042.fnal.gov</ProbeName>
  <SiteName>FNAL_PUBLIC_DCACHE</SiteName>
  <Grid>OSG-ITB</Grid>
</StorageElementRecord>
```

Hadoop Storage Probe

8

- Gets information from HDFS
 - `hadoop fs <args>`
- Runs as a cron jobs
- Reports information about:
 - Storage Element
 - Quota
 - Directories
 - Pool
- Example (from StorageElementRecord Table)

UniqueID	Measurement Type	Storage Type	Total Space	Free Space	Used Space	File Count Limit	File Count
Nebraska-Hadoop:Quota:/user/engage	Logical	disk	2,199,023,255,552	641,929,880,364	519,031,125,078	1,000,000	5,619
Nebraska-Hadoop:Pool:node15	Raw	disk	2,756,552,964,096	1,030,533,974,002	1,726,018,990,094	null	null

Xrootd Storage Probe

9

- Gets information from xrootd data servers (xrootd should be configured accordingly to enable monitoring and reporting)
- Runs as a daemon process
- Reports information about:
 - SE
 - Pools
 - Area Tokens
- Example (StorageElement Table):

UniqueID	ParentID	Space Type	Implementation	Version	Status
USCMS Xrootd:SE:USCMS Xrootd		SE	xrootd	1.1-4	production
USCMS Xrootd:Area:Xrootd Area Tokens	USCMS Xrootd:SE:USCMS Xrootd	Area	xrootd	1.1-4	production
USCMS Xrootd:Pool:cithep168.ultralight.org	USCMS Xrootd:SE:USCMS Xrootd	Pool	xrootd	1.1-4	production

Summary

10

- Storage Probes collect and report information that describes storage topology, logical and physical utilization of storage
- StorageElement and StorageElementRecord could be extended/modified when OGF standards become available
- Storage Probes are installed on a handful of the OSG Sites
- We are interested in participating in standards discussions
- We can exchange storage usage data with other accounting systems based on OGF standards